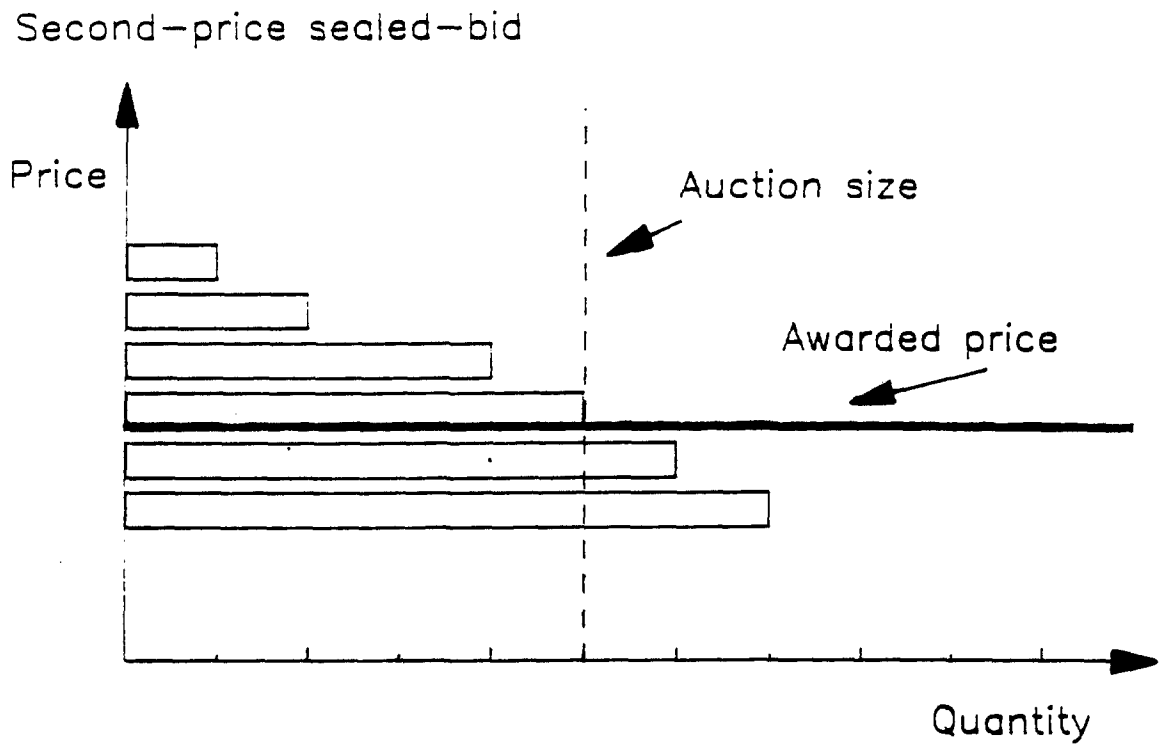
Second-price sealed-bid
Price
Auction size
Awarded price
Quantity

Figure 2

business directly by bidding at the auction rather than going through a dealer.

Descending-price open-outcry auction. This procedure is used to auction flowers in the Netherlands, hence it is referred to by academics as a Dutch auction. Bidders congregate in one room, or its electronic equivalent, and wait as the auctioneer calls out a sequence of decreasing prices. In an auction of one unit of a good or security, as depicted in Figure 3, the auction stops when one bidder is willing to pay the price called out. For multiple units, the eager bidder would be awarded the security and the auction would continue, with the auctioneer selling the remaining securities at progressively lower prices. Importantly, the strategic decision is identical to that of the first-price sealed-bid auction: The optimal bidder does not want to be too aggressive and stop the auction well above the likely market consensus, but rather, will shade his or her bid to avoid the winner's curse.[8] As a result, investors would have the same incentive to pool bids, placing customer orders at dealers.

Ascending-price open-outcry auction. The auctioneer could just as well cry out an ascending sequence of prices to the gathered bidders, stopping the auction when just enough are willing to take down the total issue. Such a price sequence is plotted in Figure 4 for the auction of a single good or security.[9] The auction of

8. This strategic equivalence was first noted by Vickrey in 1961, op. cit. In other words, what market participants refer to as an English auction is strategically identical to what academics refer to as a Dutch auction.

9. In keeping with the mirror imaging, academics term this an English auction. Indeed, in the private-values model (which we do not analyze), another equivalence proposition holds: What market participants refer to as a Dutch auction is strategically identical to what academics refer to as an English auction. Unless, of course, there is a time limit on bidding, when it is called a Scotch auction.

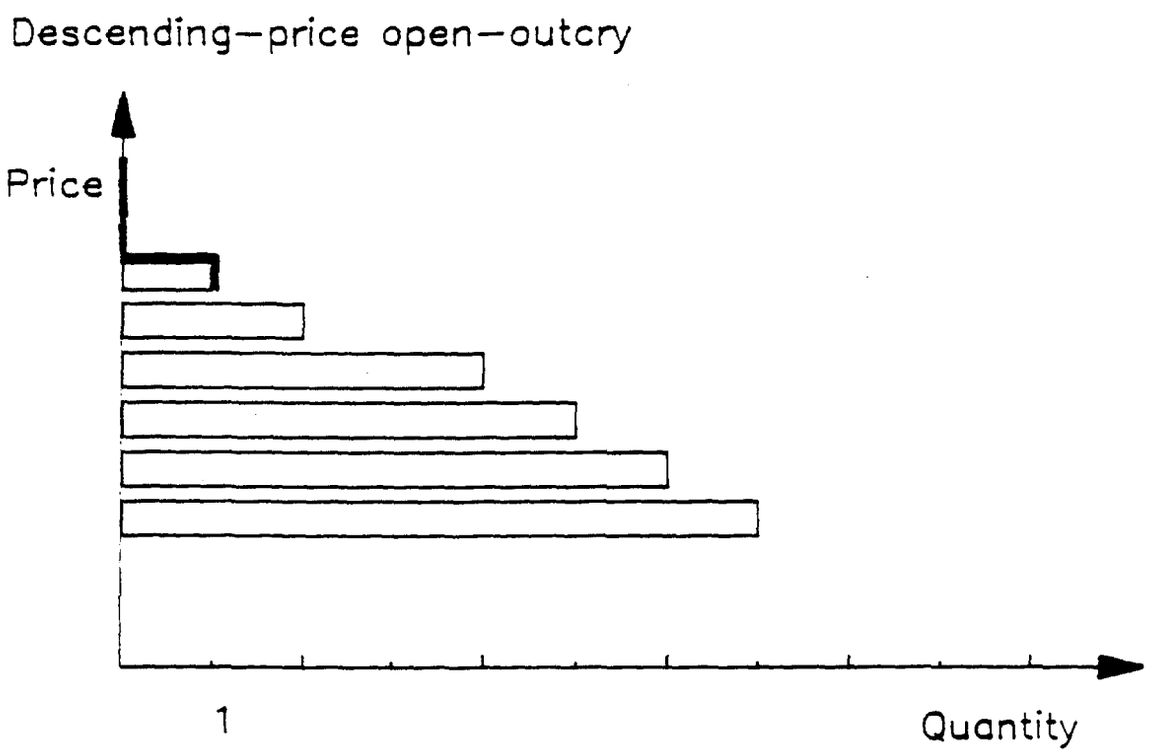


Figure 3

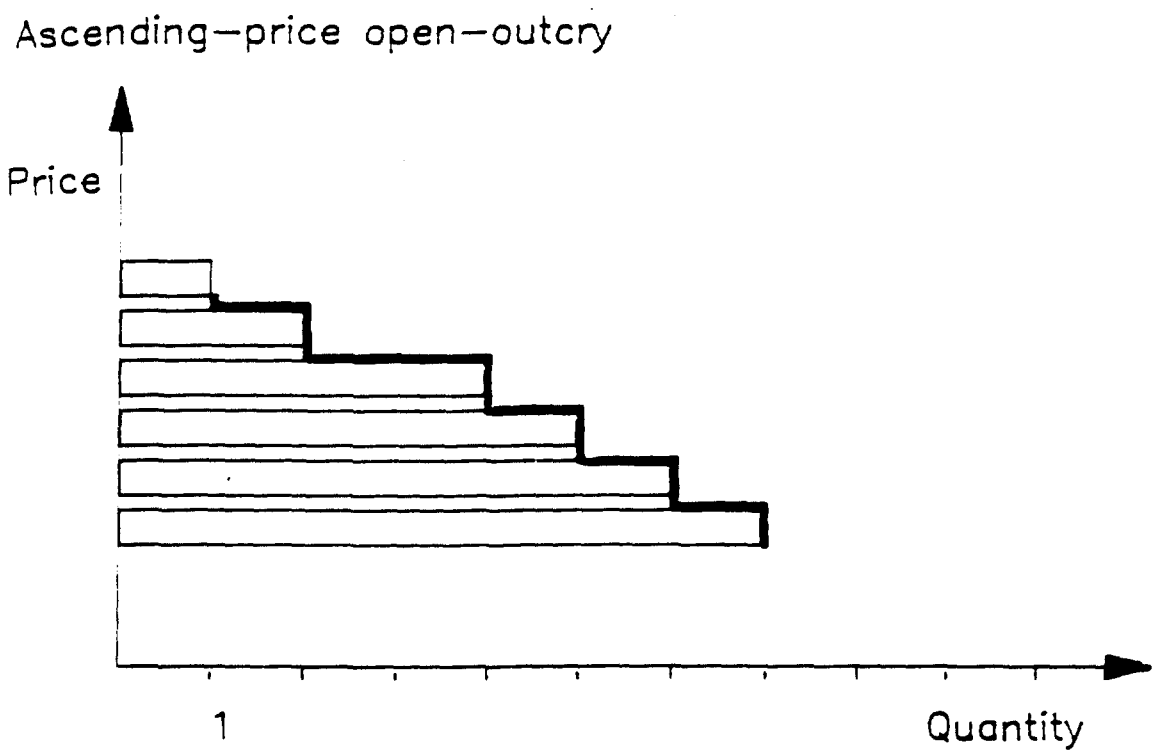


Figure 4

multiple units of a security would begin as a price was called out and all interested parties submitted their quantities demanded. The volume of bids at that price would be announced and, in successive rounds, the price would be raised until the volume demanded was smaller than the size of the issuance. When that point was reached, the seller would know that the price that was just previously called out was the highest price consistent with placing the entire issue--it clears the primary market. Everyone who bid at the top price plus some fraction of the bidders at the previous price not in the top group would receive awards at that lower price. [10] From the viewpoint of an investor, this increasing sequence of prices lessens the winner's curse, as the public outcry of bids provides information about the security's common value--namely that others also value the security highly. Besides, if an investor is truly alone in valuing the security highly, the auction stops before the price is pushed too far up when the other bidders drop out.

## III. The Information Revealed by Auction Format

In 1961, William Vickrey established that the four major auction formats provide equal proceeds to the seller. Unfortunately, this revenue equivalence requires that individual valuations are independent, or that the subjective worth of the single item on the block is unrelated across bidders. Obviously, the Treasury market violates this assumption, as the value agents place on the security reflects an imperfect estimate of the price in subsequent market trading--bidders in a Treasury auction care about the common value of the security. Against this backdrop, knowledge of other bids could importantly influence a bidder's opinion about the item.

---

10. Those partial awards could go to those who were electronically timed as placing the earliest bids or on a pro rata basis.

The expected profit from winning the auction for bidder 1 can be written:

$$(v_1 - b_1) \cdot \Pr\{b_1 > b_i, \text{ for all other } i\},$$

or the expected value of the security in secondary market trading, $v_1$, less the awarded price, $b_1$, times the probability of winning the auction, $\Pr\{\cdot\}$. Auction format determines how the bid price affects the profit from acquiring the security and the probability of winning, as well as the information revealed about the value of the security by the auction process. We first make strong assumptions about information, next discuss the strategy of maximizing the gain from participating in an auction, and last examine the information revealed during the auction.

Market Knowledge. To make matters concrete, suppose that there are two values for the security in secondary trading at the end of the day, high (H) and low (L). Further assume that bidders survey likely secondary market customers to arrive at noisy estimates of the value of the security at day's end. Since a signal is noisy, that reading on market sentiment does not reveal the true value of the security, but it does tell which is more likely. For example as Figure 5 shows, for a potential bidder, a reading of v could be associated with either a high or low price in secondary trading, but it is more likely to be associated with the high price. A market reading of v rarely occurs when the true price is L but is more frequent when the true price is H. Repeated sampling--collecting more information--would make that estimate more precise. Market contacts or pooling quotes with a dealer are ways of refining that estimate, but we assume that a bidder samples only once.

Strategy in Bidding. As a starting point, assume that each bidder's valuation is fixed (the private-values assumption), allowing

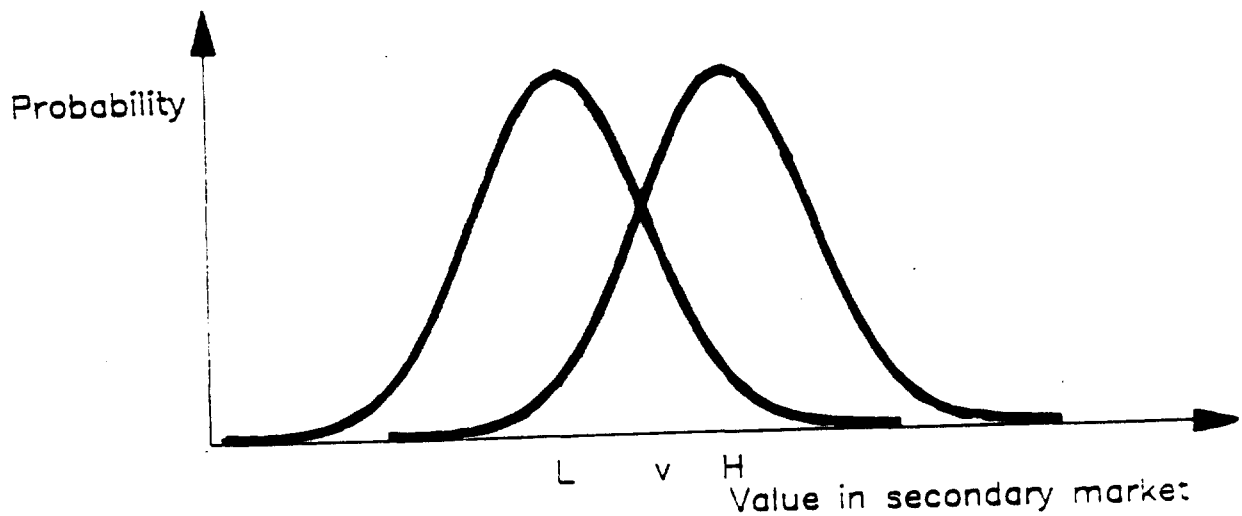


Figure 5

us to isolate the effects of the bid on the potential gain from winning the auction. It turns out that the auctions split into two groups. In a descending-price auction, the bidder must decide when to stop the proceedings. A high bid lowers the profit from victory, while making the probability of winning high. The strategic bidder trades between the two, lowering the bid relative to valuation to profit more from winning and accepting the risk of lowering the probability of winning. The decision is identical in a first-price sealed-bid auction. The winner will pay the price he or she bid and so is willing to trade probability for profit. In both auctions, the optimal bid is shaded below true valuation. But how low can a bidder go? The participant with the highest valuation only needs to beat the second-highest valuation to be assured victory. So, conditioned on a given market reading, the optimal bid is shaded to just above the expected value of the next-highest reading.

The two other auctions sever the gain in winning from the probability of winning. In a second-price auction, the winner does not pay his bid, only the second-best bid. Hence, the bidder does not control the first term, which becomes $(v_1 - b_2)$, and can set the bid to maximize the second term, the probability of winning. Consistent behavior across all bidders will make bids reflect true valuation--there is no shading. With everyone acting that way, the award will be made at the second-highest value. On the other hand, during an ascending-price open-outcry auction, there is no uncertainty about whether a bid will win. If the competition matches a bid, the auction continues. If not, there is a winner. As long as the current quote is less than or equal to the internal valuation of the security, then the bidder would stay in the race. Bids ultimately reflect true valuation. When all other bidders fall away, the winner who has the

highest valuation receives the award at just above the second-highest valuation.

It is easy to see the genesis of Vickrey's revenue equivalence theorem for the private-values model. The optimal degree of shading in descending-price and second-price auctions requires that the bidder with the highest valuation places a bid at the expected second-highest valuation. But that is exactly the outcome of the other two auction types. Hence, revenue is equal across auctions. However, this only follows because the internal valuation was held fixed. In the common-values case, that valuation depends importantly on the auction format.

The News in Auctions. A bidder receives additional information on the value of the security from the auction format. This is most obvious in a descending-price open-outcry auction, where the auctioneer calls out a series of decreasing prices until one bidder claims the item. During the auction, bidders receive negative news, as the initial quotes reveal that no one values the security in the extreme region of the price distribution; in effect, the auctioneer slices away an increasing part of the probability area, showing that the true valuation must lie in the lower price region (Figure 6). Thus, this information leads a strategic bidder to reduce the valuation of the security below his or her initial reading. Sealed-bid auctions provide this same negative news. A strategic bidder calculates the expected profit from winning, which implies that his or her valuation must be higher than anyone else. Hence, planning on winning requires trimming the expected valuation of the security below the initial reading.

Only an ascending-price auction differs. As the auctioneer calls out an increasing price list, bidders receive news that

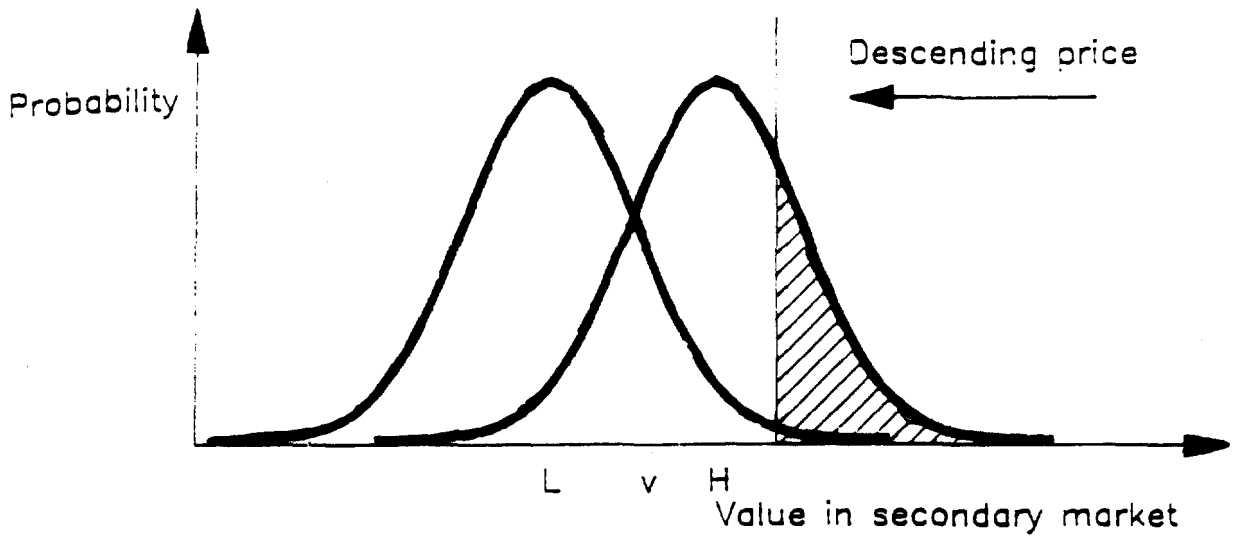
Probability
Descending price
L
v
H
Value in secondary market

Figure 6

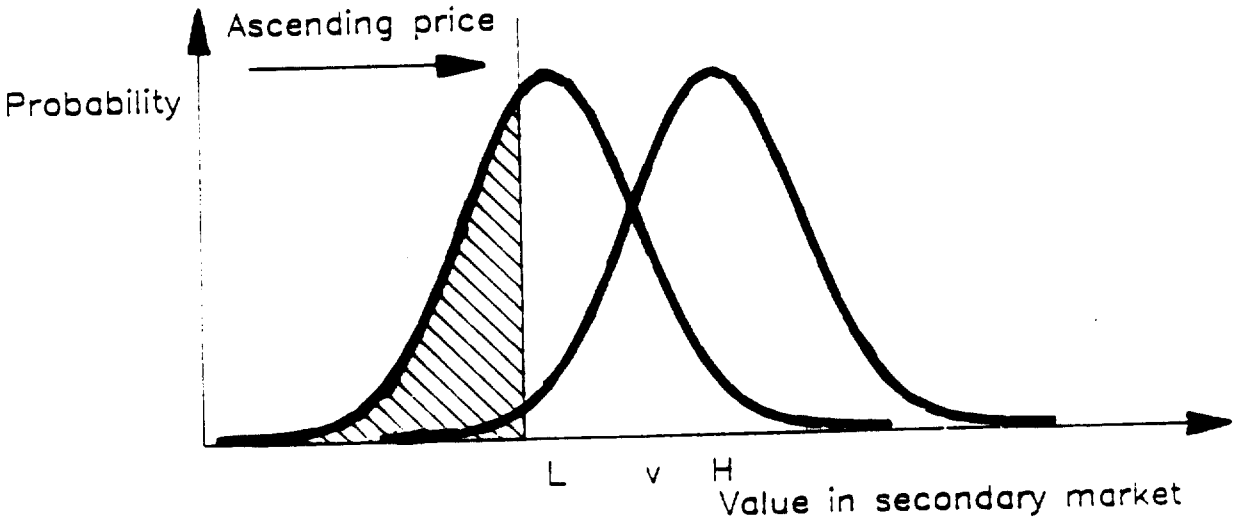
Ascending price
Probability
L
v
H
Value in secondary market

Figure 7

participants prize the security more highly than those low quotes. In effect, the auctioneer trims the low end from the price distribution, revealing that the true market values is likely higher (Figure 7). Thus, an ascending-price open-outcry auction produces the highest expected value of the security. With bidders conditioning their behavior on the highest expected value of the security and limiting their bid shading, revenue is greatest for this auction format.[11]

## IV. The Potential for Illicit Profit in Auctions

As already shown, the current auction setup elicits one form of strategic behavior: Because awards are priced at the bid, the rational participant shades a bid to avoid the winner's curse. As a result, customers pool their bids with dealers so that a combination of bids can, by a law of large numbers, be appropriately cast. However, the auction form breeds two other types of strategic behavior. First, in single-dealer cornering, a dealer may combine with a customer to corner a significant portion of one auction--70 percent if it does not break the rules. Alternatively, in collusive combining, a group of dealers can conspire to accomplish the same end. In a sealed-bid auction, either ring need only make a slightly more aggressive bid than the other participants expect to garner the lion's share of awards. Indeed, the pricing mechanics of Friedman's proposal make these cheaper propositions. The pool could corner the issue by bidding substantially more than the market consensus but pay a price closer to the mass of the distribution that marks the other bids.

Clearly, single-dealer cornering and collusive combining are similar--a ring has the same intent on squeezing the market as an

---

11. This was shown formally in 1982 by Milgrom and Weber, op. cit., in Theorem 11.

individual firm. However, the informational requirements and incentives for these two type of strategic behavior vary across auction type and actions taken to combat one might make the other more likely.[12] This section will analyze the collusive potential in auctions, but first, it is important to understand the incentive behind cornering an auction--or, how one variety of a squeeze can work.

How a Corner Works. Salomon's misdeeds were directed toward paying more for a larger quantity than anyone else at an auction. In fact, the potential for profit in a squeeze lies in the interaction of the three main trading fora for Treasury securities: the when-issued market, the Treasury auction, and the secondary market. Those markets are represented by the three panels of Figure 8, arrayed by time--before, at, and after the auction. As the right panel shows, the price of a Treasury security must satisfy the ultimate holders of securities (pension funds, insurance companies, mutual funds, and the general investing public), seen as the intersection of their downwardly sloped demand schedule with the vertical Treasury supply schedule.[13]

Current auction procedures, however, get securities to those holders indirectly, through the intermediation of dealers; as shown in the middle panel, the demand derived from current and anticipated customer orders produces a flatter and more inward schedule at the

12. The manipulative possibilities open to a single trader with large capital are discussed in Gary L. Gastineau and Robert A. Jarrow, "Large-Trader Impact and Market Regulation," *Financial Analysts Journal*, (July/August 1991), pp. 40-51.

13. Demand schedules are labeled according to time: WD, during when-issued trading; AD, at the auction; and SD, during trading in the secondary market.

The three main markets for Treasury securities

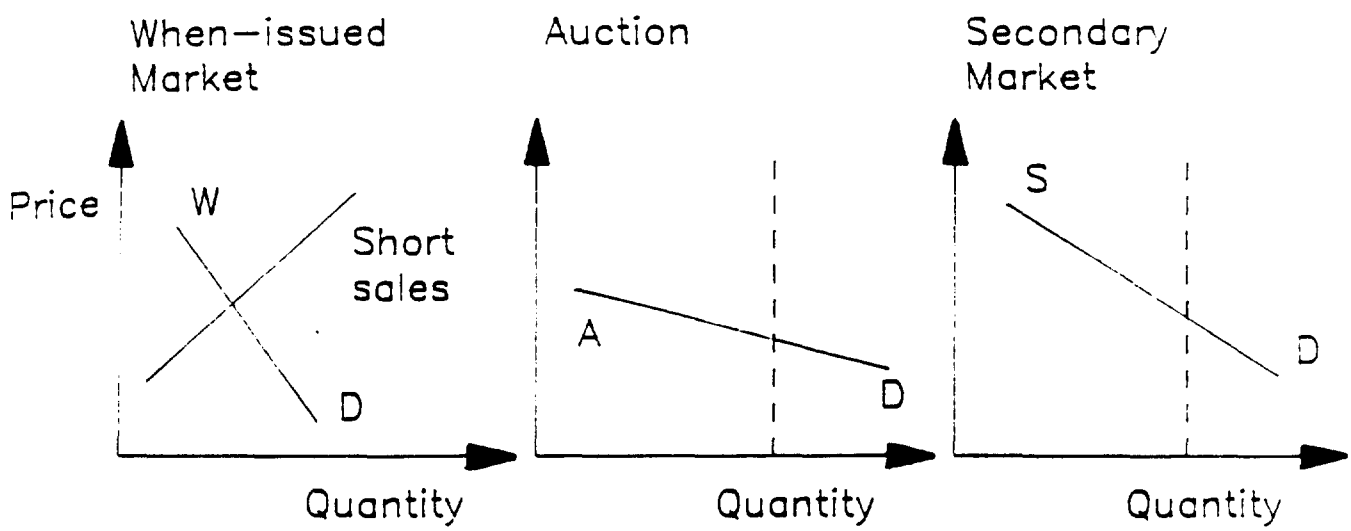


Figure 8

auction--which is the optimal shading of bids in response to the winner's curse.

Additionally, an investor can purchase the security before the auction, as long as he or she can find someone willing to sell it short. The when-issued market, shown in the left panel, matches those parties. Those seeking secure ownership rights trace a downwardly sloped demand schedule, while those willing to sell what they do not have yet make up the short-sale schedule. Selling a security before the auction involves a quantity risk, as short sellers may not win awards at the auction to cover their open positions and, hence, will have to borrow the security after the auction settles to make delivery. Accordingly, the when-issued price should clear above the expected auction price.

The next set of diagrams (Figure 9) depicts the cornering of an auction. Short sales are made at a price just enough above the anticipated auction price to indemnify the sellers of the likely risk at the auction. Those sellers, however, turn out to be wrong about the auction, for while the market consensus coalesces around bids consistent with the demand schedule AD in the middle panel, one party nips in with bids that shift the actual demand schedule to A'D'. The cornerer exploits the sealed-bid nature of the auction: By bettering the market consensus, the schemer wins the bulk of the awards (measured by the horizontal distance between the two demand schedules).[14]

14. A manipulator could bypass the auction by amassing a controlling position in either when-issued or secondary market trading; to affect that strategy, purchase orders would have to be spread across many sellers in an effort to hide the intent to corner from the general market.

How a corner works

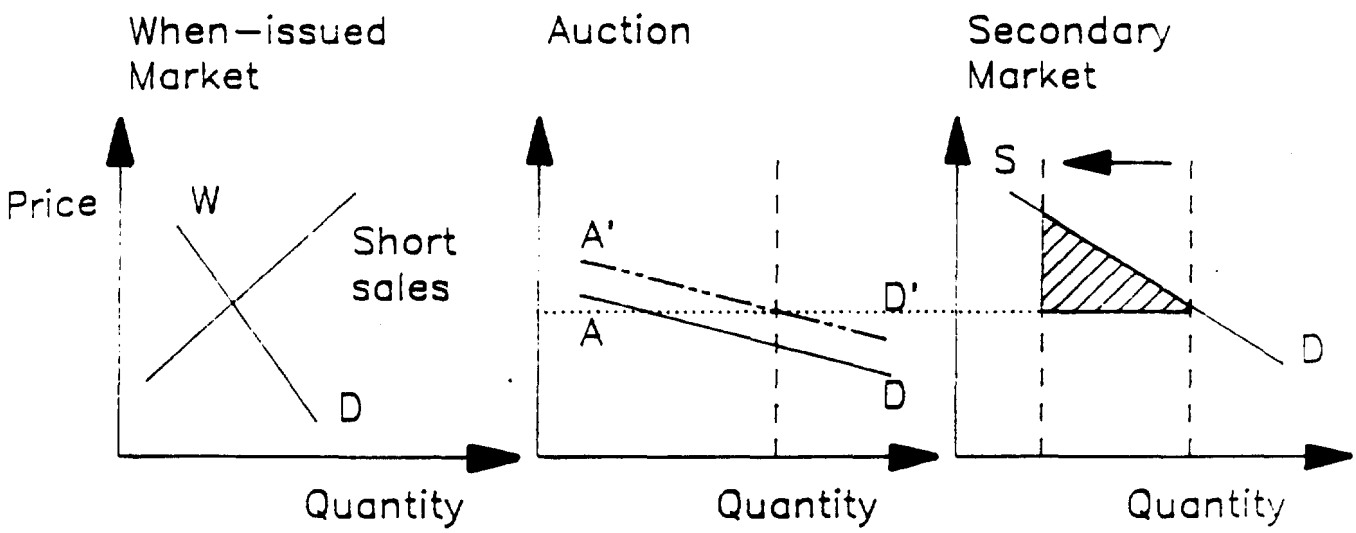


Figure 9

Since other parties cannot react, the Treasury receives only a modestly higher price for its auctioned securities, but the major price action awaits secondary market trading. The cornerer restricts the supply of the security in the secondary market (seen as the inward shift in the vertical supply schedule in the right panel), so that the price that clears that market is well above the auction price. From there, the cornerer slowly unwinds that position, expanding market supply to sell at prices above the ultimate level determined by the buy-and-hold ownership of Treasuries. In effect, the cornerer acts as a discriminating monopolist, carefully regulating sales to push the price down to earn all the revenue given by the area under the demand schedule. The cornerer's cost, on the other hand, is given by the unshaded rectangle, leading to the profit given by the shaded area.

Indeed, the shaded area understates the potential for profit from a market squeeze. While the issue remains in cornerer's control during secondary trading, short sellers are obliged to borrow the security to make delivery. That transaction is one side of a repurchase agreement in which the owner of the desirable security--the cornerer--lends it to a short seller in return for cash at a preferential borrowing rate. In effect, by creating a demand for the issue, the cornerer can finance his or her position at a below-market borrowing rate, increasing the total profit from the squeeze.

The when-issued market plays two important roles. First, early trading allows the market consensus to coalesce quickly, providing a usually accurate forecast of the auction price. That market serves importantly in the "price discovery" of the appropriate rate on the to-be-auctioned security, which tightens the spread of bids and allows the cornerer to bid only slightly higher than that to be assured awards. Second, a group of thwarted bidders--those who

shorted in the when-issued market--are forced to the secondary market to close their position. Their surprising presence makes the demand schedule less price sensitive, as there is no substitute for the security that they promised to deliver. As a result, as long as they keep open their positions, short sellers will need to borrow the desirable security, providing the cornerer favorable financing in the RP market.

Thus, the successful cornerer preys on three attributes of the current setup:

- When-issued trading creates a core of reliable demanders for the auctioned security (those who earlier sold short);
- The first-price method of allocating awards reduces demand at the auction and makes that demand more price sensitive; and,
- Sealed bids allow a cornerer to place bids only marginally better than the consensus and win all the awards.

Thus, current procedures promise profit in cornering a Treasury auction.

The Potential for Collusion. Clearly, one dealer with adequate capital could take advantage. A harder question to assess is if the setup of an auction may entice a group of dealers to conspire together in the attempt.[15] The theoretical analysis of the incentives for collusion in auctions proceeds as follows: Suppose, for argument's sake, that a small set of dealers willfully plan together to purchase all that is sold at an auction, intent on extracting profit from those not in the ring. Together, they agree on

15. Some of the analytic difficulties in modelling the strategic interaction of a group of potentially cooperating bidders can be seen in Daniel A. Graham and Robert C. Marshall, "Collusive Bidder Behavior at Single-Object Second-Price and English Auctions," *Journal of Political Economy* vol. 95, 1987, pp. 1217-1239.

a price just above the market mean that is sure to win all the awards. However, a sealed-bid auction tempts each of the co-conspirators to move just above the agreed-upon price and steal awards--the cartel likely will not hold.[16] Hence, incentives in the classic first-price sealed-bid auction are structured so as to make collusion unlikely. On the other hand, in an ascending-price open-outcry auction, such a cheater amongst the cheaters has to show his or her hand, making such cheating less likely.[17] The cartel will hold.

By this theoretical argument, one might surmise that the Treasury's first-price sealed-bid auction, at least, protects against the willful joining of dealers to exploit the Treasury and other dealers. Unfortunately, there is a gap between models and reality, as the rule limiting awards to 35 percent of the issue paradoxically turns incentives back toward collusion. If a conniver plays within the lines of the 35 percent rule, a cheater on his fellow cheaters will not win enough securities at the auction to control the secondary market. Hence, tough enforcement of quantity limits binds conspirators closer together.

More to the point, theoretical analyses of collusion assume that a small number of colluding parties share information, which neglects the multiple arenas in which dealers compete.[18] Dealers

16. This would also hold for a descending-price open-outcry auction. The first escapee from the pool stops the auction before the others can react.

17. Even if bidding is secret, the other members of the cartel will know by the price movement that someone cheated.

18. As Marc S. Robinson writes of his proofs of the incentive structure of auctions: "An important condition of the above theorems is that no private information remains inside the cartel. If bidders have private information, they will frequently have positive expected

(Footnote continues on next page)

will not cooperate in auctions if it jeopardizes their secondary market trading.[19] Given the large number of participants and the apparent mistrust among dealers, auction form is unlikely to bring them together. Thus, from a public policy standpoint, the chief concern should be combatting manipulative actions by a single dealer--the rogue with capital--that call into question the integrity of the market.

## V. A Closer Look at the Friedman Proposal

In the editorial page of the *Wall Street Journal*, Milton Friedman repeated a proposal he advanced in 1959 concerning the best way to auction Treasury securities.[20] Essentially, Friedman argues for a second-price sealed-bid auction. In the one alteration to current practice, the Treasury would give up price discrimination, awarding securities at a uniform price rather than the bid price. Friedman asserts that the switch would end cornering attempts by eliminating the profit potential in market manipulation. And sounding

---

(Footnote continued from previous page)
profits, even if they do not cooperate. For a stable cartel to be formed, every member must prefer both participation and compliance." Quoted from "Collusion and the Choice of Auction," *The Rand Journal of Economics* vol. 16, Spring 1985, pp. 141-145.

19. The existence of inter-dealer brokers is one sign of the level of mistrust among dealers. These intermediaries provide anonymity to dealers in transactions between dealers, who are reluctant to phone their competition directly and show which side of the market they are on.

20. "How to Sell Government Securities," *Wall Street Journal*, August 28, 1991. Merton Miller also has embraced this reform, as quoted in Diana B. Henrique, "Treasury's Troubled Auctions," *New York Times*, September 15, 1991.

paradoxical. Friedman argues that total revenue would be higher by surrendering the ability to discriminate across price bids.

The Consequences for Cornering. As seen in the previous section, the current Treasury technique reduces demand at auctions, as well as making it more price sensitive, relative to that determined by the buy-and-hold ownership of the long-time investor. This is the rational response to the Treasury's discriminating pricing: The investor shows less of his true consumer surplus to a seller whose stated intention is to seize it. But with this difference in demand, a cornerer can buy at the auction at a price not much influenced by his own demand and sell in the secondary market to a larger group of investors.

Moving to a common-price setup permits demand at the auction to reflect the true nature of investor preferences. With no frictions, investors can bypass the dealer intermediaries and directly bid, sharing the resulting savings with the Treasury. As a limiting case and in terms of our three-figure determination of Treasury prices, second-price awards would make the auction demand curve identical to the secondary market demand curve (Figure 10). Against this backdrop, the cornerer of an auction would place surprising bids that shift the demand schedule from AD to A'D'. The horizontal distance of that shift represents the cornerer's awards, or the extent that secondary market supply can be restricted. But, as seen in the right panel, the investors who were unwilling to pay the auction price will be unwilling to pay the secondary market price. Now, the cornerer acting as a discriminating monopolist minimizes loss, seen as the shaded triangle, rather than maximizes profit. Clearly, as the Hunt brothers learned to their regret, you cannot profit from cornering a market with invariant demand, because you ultimately must

y

s, as
ed by
e

The effect of Friedman's proposal on cornering

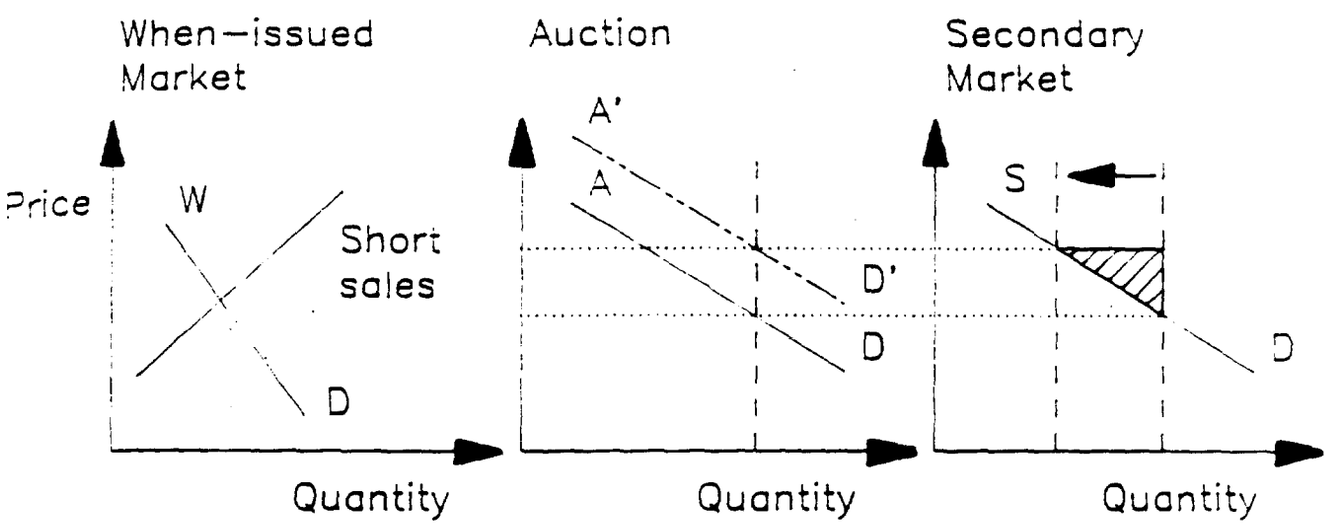


Figure 10

e
and,
by
of

tion

ectly
ing

ıst
ɔids

xtent
the
price

en as
ıe

nust

sell the security to those from whom you bid it away. Thus, under Friedman's assumptions, cornering would be eliminated by removing the potential for profit.

This result, however, requires that the switch in auction technique completely unifies the primary and secondary markets. Even after the adoption of common-price awards, presence at auctions still may be limited to a segment of the investor populace, perhaps to those who are more price sensitive. Those who sold short in the when-issued market quickly want to cover their positions at the auction. Also, participants at an auction face uncertain outcomes, since they may not be awarded securities if they have not appropriately cast their bids. Those particularly adverse to this quantity risk well may delay purchase to secondary trading. Most importantly, direct bidding requires incurring the fixed costs of assuring payment and arranging for the placement of bids--the prospects for which depend on the pace of automation and regulation. As a result, the infrequent purchaser may remain in the secondary market. In other words, Friedman assumes that dealers exist solely to shade bids because of the Treasury's discriminatory pricing. If, however, dealers provide any other service in the distribution of securities, then a wedge remains between the auction and secondary market demand schedules. A sufficiently large wedge represents an opportunity for manipulation.

The Consequences for Revenue. The algebra required to calculate an optimal plan in a multiple-unit auction pyramids into incoherence. No analyst yet has worked through the strategic implications of a large core of bidders carving up a block of securities. The logic of the single-unit case, however, suggests that the extent of bid shading can be extreme. In a first-price auction of multiple units, a strategic bidder does not have to beat the

participant with the next highest valuation to win but only must better the middle of the pack of bidders.

If we step away from the explicit modelling of bidder behavior, the revenue implications can be spelled out in terms of shifts in the demand schedule for the auctioned security, as detailed in Henry Goldstein's 1962 analysis.[21] As Figure 11 shows (which repeats the middle panel of our three-figure determination of market prices), part of the Treasury's total revenue owes to its charging winners the price that they bid, which for the current practice is measured by the area under the demand schedule labeled "first-price." That price discrimination, however, discourages some demand, as investors shade their bids for fear of the winner's curse. Adopting Friedman's second-price system turns part of that surplus back to the bidders, shifting out the demand schedule to that labeled "second-price." Under a first-price scheme, the Treasury would have to work down the inner demand schedule, awarding securities at lower prices to place the total issue (marked by the vertical dashed line). Under the second-price scheme, one price, depicted by the horizontal line, exhausts the issue. The consequences for revenue depend on whether the area of the first triangle, the loss from the inability to price discriminate, outweighs the area of the second triangle, the gain from added demand.

Current Treasury practice seizes all the consumer surplus--the entire area under the demand schedule from the maximum price that a bidder shows to the stop-out price. However, the logic of bid shading suggests that rather than observing--and exploiting--the entire upper part of the demand schedule, the Treasury is confronted

21. "The Friedman Proposal for Auctioning Treasury Bills," *Journal of Political Economy* Vol. 70, August 1962, pp. 386-392.

The effect of Friedman's proposal on revenue

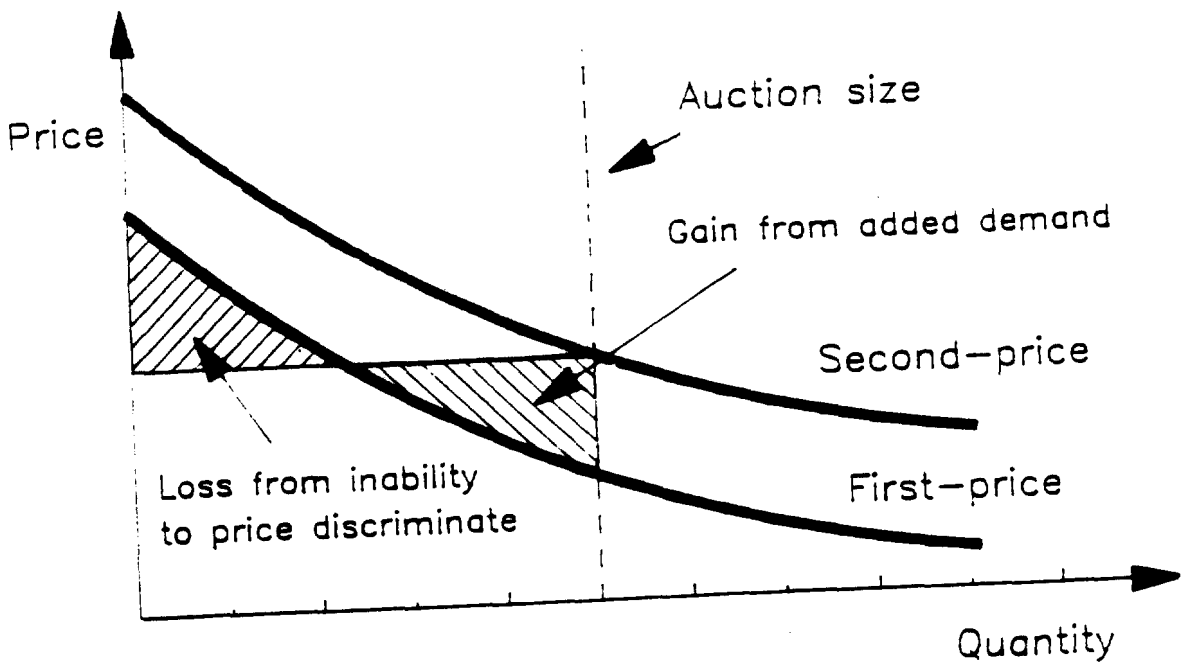


Figure 11